Analysis of User Behavior and Traffic Pattern in a Large-Scale 802.11a/b Network

Jaehwan Lee[+], Sunghyun Choi[*], Hanwook Jung[+]

[+]Convergence Laboratory
Korea Telecom
tasmea@kt.co.kr, hanuk@kt.co.kr

[*]School of Electrical Engineering
Seoul National University
schoi@snu.ac.kr

Abstract- Today, IEEE 802.11 wireless LAN (WLAN)-based networks are typically used at large-scale conferences or meetings. The 59th IETF standard meeting was held in Seoul, February/March 2004. During the meeting, an 802.11a/b dual-mode WLAN network was provided for participants, and we measured users' usage logs via Simple Network Management Protocol (SNMP) and System Logs (SYSLOGs). This measurement covers 1293 users and 32 APs for 5 days. In this paper, we analyze the network statistics and user behaviors from the measurement results. One major distinction of this work from other existing literature is that we consider co-located 802.11a and 11b networks, and evaluate their network characteristics in a comparative manner. We find that (a) user arrival and traffic patterns highly depend on the meeting schedule, specially in case of APs located in meeting rooms, (b) 11a traffic per user is more intensive than 11b traffic, (c) sojourn time spending with a single AP is similar between 11a and 11b, but 11b client hands off more frequently, and (d) the 11a hand-offs occur only between relatively closer APs compared to the 11b's case. The measurement analysis in this paper should be useful for designing and deploying large-scale conference-style WLAN networks, especially, using 802.11a/b dual-mode APs.

Keywords – IEEE 802.11, Traffic Pattern, User Arrival Pattern, User Mobility, Hand-off Asymmetry

I. INTRODUCTION

Today, IEEE 802.11-based wireless LAN (WLAN) is emerging as the major means to provide networking connectivity wirelessly. Even though the number of the WLAN users increases drastically today, it is our understanding that many issues related to how to deploy, design, and maintain the WLAN networks still deserves more research and engineering efforts.

Currently, the most popular form of WLANs is IEEE 802.11b, which provides about up to 5 Mbits/s (Mbps) network throughput based on up to 11 Mbps data transmission rates at 2.4 GHz unlicensed bands. One stimulus for expanding WLAN usages is the launch of higher-speed WLAN-based services. IEEE 802.11a protocol provides up to about 24 Mbps network throughput using up to 54 Mbps data transmission rates. In addition, 5GHz frequency bands used for 11a are relatively broader and cleaner (i.e., with relatively less interference) compared to the 2.4 GHz bands used for the 802.11b WLAN. We expect that 11a and/or its variant will be one of major WLAN interfaces in the future.

The 59th Internet Engineering Task Force (IETF) meeting was held in Seoul, February/March 2004. Korea Telecom (KT) offered the network connections for the meeting, including wireless access based on both 802.11a and 11b. During this meeting, the network traffic was measured, involving 1293 users and 32 Access Points (APs) for 5 days. This traffic amount and the number of users are significantly larger than those reported in the existing literature [1][2][3][7].

We in this paper present an analysis of user behavior and traffic pattern in a large-scale WLAN based on our measurement. To our best knowledge, all the existing WLAN measurement reports have been made for the currently popular 11b WLANs. On the other hand, our measurement analysis is from an 802.11a/b dual-mode WLAN.

The rest of the paper is organized as follows. Section II briefly introduces other related work in comparison with ours. The 59th IETF wireless network environment is presented in Section III. Section IV presents the network statistics and analysis such as user arrival, traffic, session, and user mobility patterns. Finally, the paper concludes with the summary of our findings in Section V.

II. RELATED WORK

There have been various studies related to the WLAN measurement and its analysis. One of the earliest reports is by Tang and Baker [1]. It analyzes the network traffic from 74 users for 12 weeks. This research mainly focuses on the users' behavior pattern during the usage of the wireless network. With the current trend of large-scale WLAN deployment at universities, researchers at Dartmouth College present significantly large and broad trace of campus-wide 802.11 network [2]. This research uses 476 APs covering over 161 buildings and analyzes network traffic, user mobility, client card and AP activities, protocol usage pattern by gathering Simple Network Management Protocol (SNMP), SYSLOG, and sniffed packet headers. While this report from Dartmouth covers overall network and user mobility patterns, North Carolina's researchers focus on user association and mobility pattern using their network measurement [3].

The closest research to our work is by Balachandran *et al.* based on the measurements during ACM SIGCOMM'01 conference [7]. While other reports focused on campus or public area networks, this paper analyzes a conference type network, where many clients move frequently and rush to the meeting rooms based on the conference schedule. They gather the information of 195 users from 4 APs for 2.5 days. Our research expands this measurement on a larger network and a broader population; we traced 1293 clients using 32 APs for 5 days. The most distinctive feature of our work is the measurement of 802.11a network, which is co-located with the 802.11b counterpart. We deployed 11a/b dual-mode APs for this meeting so that we could get the detailed information about 11a user's behavior pattern and network usage compared to the 11b users. Even though the number of 11a users is 64, which is relatively very small compared to 11b's, the general pattern of the 11a network is very useful because 11a/b dual-mode APs are expected to be more popular in the near future.

III. NETWORK ARCHITECTURE AND MEASUREMENT METHODOLOGY

A. *Network Architecture*

The 59th IETF meeting was held for five days from February 29th to March 4th, 2004. Sessions were held at many large and/or small rooms simultaneously. On the first day, there were only introductory training sessions for new comers. Beginning the second day, i.e., from March 1st to March 4th, the morning sessions stared at 9:00 am, and the evening sessions finished at 10:00 pm (except March 2nd). The detailed meeting schedule can be found at http://ietf.org/meetings/agenda_59.html.

The wireless network covered conference rooms, lobbies, and other miscellaneous areas. We used the Cisco AP products, namely, 1200AP. These APs have 11a/b dual-mode wireless interface. Cisco 1200AP supports the 802.11g mode as well, but we had no experience with 802.11g operation, and hence we decided to serve only 11b network at the 2.4GHz band. We used number 1, 6, and 11 channels at 2.4 GHz, since client network adapters from the U.S. and some other countries support only channels from 1 to 11, unlike cards from Korea supporting channels from 1 to 13 [8]. At the 5 GHz bands, we used 8 channels from 5.15 to 5.35 GHz.

As the general public WLAN network, clients get their IP address via Dynamic Host Configuration Protocol (DHCP) after the association with an AP. All the APs were connected in the same IP subnet so that layer-3 hand-off requiring mobile IP support is not needed.

B. *Measurement Methodology*

To gather statistics, we employ two methods. First, we collect SYSLOG data from all the APs in operation. Cisco 1200 AP's SYSLOG contains the information related to many network operations including association, dis-association, hand-off, and other event/error messages [6]. Therefore, we could extract the session information, the number of associated clients per AP, and the user mobility pattern from the collected data. These data are separately available for 11a and 11b interfaces, and hence the 11a session pattern can be compared with that of 11b. The other way to collect the information is SNMP. We collect the SNMP data from APs every 30 minutes (from February 29th to March 2nd) or every 5 minutes (on March 3rd and 4th). The SNMP data consist of traffic information per every virtual LAN interface, and hence we can differentiate wireless network data from other data such as management packets, broadcasting packets from the wired interface, etc.

All the data were collected by one server continuously during the entire 5 meeting days. After collecting the data, we converted the raw data to the meaningful information. In the next section, we present the measurement analysis results.

IV. RESULTS AND ANALYSIS

A. *User Arrival Pattern*

We show the distribution of the number of associated clients with APs. Fig. 1 shows the overall distribution of total clients. We deployed APs at the conference rooms as well as other areas including lobbies and restaurants, and hence we divide the APs into two categories as follows:

- *Type A*: the set of APs located in large meeting rooms. These APs usage pattern will be like a conference type; many people arrive at the start of a session, and depart at the end of the session. During a break time, user traffic decreases much lower than that of the session times. The number of type-A APs is 13.
- *Type B*: the rest of APs, i.e., those not belonging to Type A. These APs' traffic is also correlated to the session time schedule, but the effect is much more minor than that of Type A. Type-B APs are located outside the large meeting rooms. The number of type-B APs is 19.

Note that the total number of unique 11b clients is 1229, and that of 11a clients is 64. Considering that total registered participants were 1408 persons, it is very likely that most participants brought their laptops during the meeting, though some people might have brought two or more WLAN adapters. Another interesting issue is the percentage of 11a adapters. Even though the 11a adapters were introduced at the U.S. retail market from early 2002, only about 5% of the users used the 11a clients. Due to Korean Government's regulation for 5 GHz, most Koreans do not have 11a adapters. If we assume that all the 64 11a clients were owned by non-Korean people, the 11a clients represent at least 8% of the entire WLAN clients from abroad.

Fig. 1 shows that most users connect the network during the session time. This user arrival pattern is consistent with the previous study result in [7]. The association arrival highly depends on the session schedule. In Table 1, we present the simple statistics about the total associated clients with the overall network. To analyze the relationship between the session time and the associated clients, we divide time into two categories, namely, session times and non-session times. According to Table 1, overall associated clients increase during session times. It is because session times are generally when people participate in the meeting actively, and non-session time includes midnight, when people are actually gone for a day. One remarkable thing is the client distribution of types A and B with respect to session times. During session times, the 11b clients associated with type-A APs is about four times of those with type-B APs. However, during non-session times, there are more clients associated with type-B APs than type-A APs. The same phenomenon is also observed for the 11a case. Therefore, the arrival patterns to type-B APs are different from those to type-A APs. This result shows that APs inside the conference rooms are heavily affected by user's arrival pattern, but APs outside rooms do not need to consider that pattern correlated with the session schedule.

From Fig. 2, we see that the type-A APs handled average 17.9 clients. The average number does not seem very high, but when the users were crowded, the number of clients associated with an AP increases over 100. It is not easy that one WLAN interface handles over 100 clients simultaneously, but fortunately the traffic was not so heavy that the AP could operate the network without difficulty.

B. *Traffic Pattern*

Uplink and downlink traffic patterns can be analyzed from the SNMP data. Usually, the amount of downlink traffic is more than that of uplink traffic because most clients receive

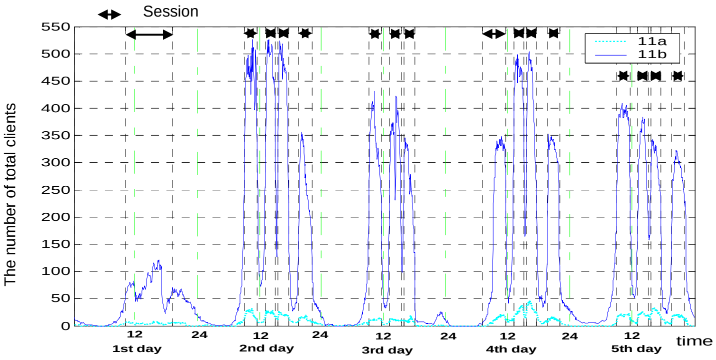

Fig. 1. The number of total clients using the wireless network

TABLE I
The number of total associated clients by AP type and time in terms of conference session or not

AP Time	All	Session	Non-session
All APs	7.59/ 4/ 47	15.42/ 15/ 47	2.74/ 1/ 25
Type A	4.22/ 0/ 39	9.97/ 10/ 39	0.65/ 0/ 18
Type B	3.37/ 3/ 13	5.44/ 5/ 13	2.09/ 1/ 10

(a) The number of 11a clients (mean/ median/ max)

AP Time	All	Session	Non-session
All APs	131.9/ 47/ 526	291.6/ 323/ 526	32.99/ 11/ 309
Type A	97.52/ 10/ 469	233.6/ 270/ 469	13.51/ 0/ 243
Type B	34.42/ 28/ 138	57.97/ 57/ 138	19.83/ 10/ 103

(b) The number of 11b clients (mean/ median/ max)

the data when using HTTP, FTP, and other applications. Fig. 3 shows the time-varying traffic pattern. It shows that the traffic pattern also depends on the meeting schedule, like user arrival pattern. The 11a traffic amount is lower than that of 11b traffic, because there are much more 11b users than 11a users. However, though 11a users represent only about 5 percent of 11b users, 11a traffic is one sixth of 11b traffic (see Tables 1 and 2). It means that 11a users consume more bandwidth than 11b users. It is because the transmission rate (and hence the throughput performance as well) of 11a is higher (i.e., up to 54 Mbps of 11a vs. up to 11 Mbps of 11b), and the number of 11a users is small so that 11a users can access the network without contending with many other 11a users. Another different point is the maximum average traffic of 11a and 11b. In Fig. 3, generally 11b traffic is heavier than 11a traffic, but the maximum average traffic of 11a is over 35 Mbps, which is approximately a double of the 11b's maximum. It is again thanks to a larger throughput of 11a.

Fig. 4 shows the relationship among the number of associated clients, total downlink traffic, and the average downlink traffic per client. All the data are from 11b interfaces associated with a certain type-A AP. The trend of this AP is similar to the overall patterns. Generally, the user arrival and traffic patterns are correlated, although they do not match exactly. From Fig. 4 (b), the traffic is very small during non-session times, and the average downlink traffic per client during session-times is about 22.3 kbps, which is actually too small compared to the 802.11b's peak transmission rate, i.e., 11 Mbps.

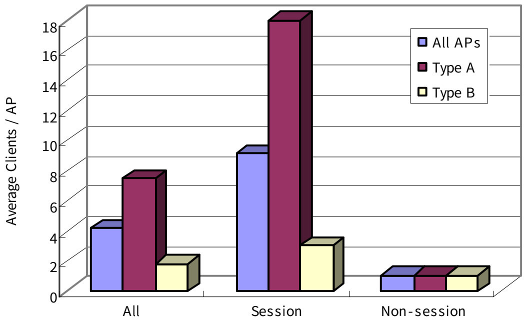

Fig. 2. The average number of associated clients per AP

C. *Session Pattern*

How long each client's session lasts as well as how many APs each client visits during the lifetime of its session are very important characteristics related to clients' usage behavior. Fig. 5 shows Cumulative Distribution Function (CDF) of the session time duration, which is defined as the time interval from the time when a certain client associates with an AP to the time when it disassociates with the same AP or another AP. In SYSLOG, we find the association and disassociation events from all APs. When a client associated with one AP switches to another AP (i.e., a hand-off situation), we interpret that the session continues because the client can use the Internet connection at the application level. Therefore, our definition is different from that of Balachandran's [7], in which the session time duration is defined by a period during which a client stays associated with an AP; we refer to this period as sojourn time as discussed further below. In the case of 11b, the average/median/maximum values of the session time duration are 2126/818/32447 seconds, respectively. On the other hand, the mean/median/maximum values in the case of 11a are 1341/318/24110 seconds, respectively. Generally, 11b clients' session time durations are longer than those of 11a clients. It is due mainly to the difference of hand-off frequencies between 11a and 11b. Fig. 6 plots the CDF of the hand-off event numbers during the lifetime of a session. The median value of hand-off number is 0. It means that over half of the entire clients did not roam to other APs during their session lifetime. However, 11b client tends to roam more frequently than 11a clients. The mean value of 11b's hand-off number is 1.92, which is more than 11a's mean value, 0.72.

Fig. 7 shows the CDF of the sojourn time, which is defined as the time period between an association with a certain AP and a disassociation with the same AP (i.e. no handoff), or movement to another AP (i.e. handoff). Therefore, if a client does not hand off, the session time is the same value as the sojourn time. When a client hands off one or more times, the session time is the sum of all sojourn times. The definition of sojourn time is conceptually the same as the Balachandran's definition of user session time [7]. Note that the sojourn time distributions of 11a and 11b are almost the same. The mean/median/maximum values of 11b's and 11a's sojourn times are 761/151/7199 and 818/186/7165 seconds, respectively. Therefore, the major reason for the difference between 11a and 11b session time durations is not the sojourn times, but the hand-off event numbers. The reason why the 11b's session time is longer than 11a's is not very obvious.

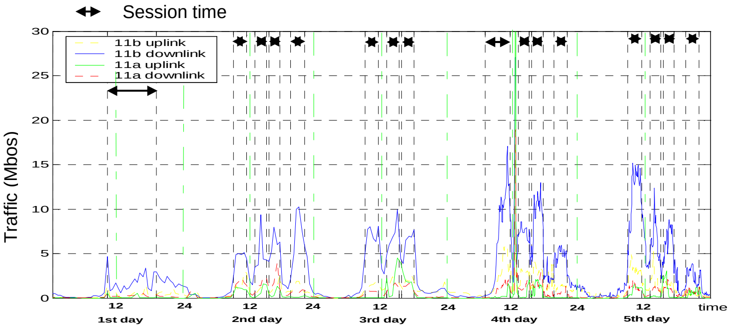

Fig. 3. The total wireless traffic pattern

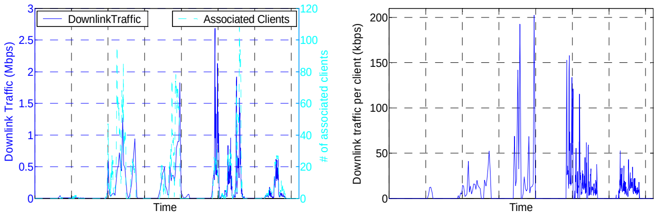

(a) The number of associated clients and downlink traffic per AP

(b) Downlink traffic per client

Fig. 4. The downlink traffic pattern of a certain type-A AP

TABLE II.
The average served traffic by AP type and time (kbps)

AP \ Time	All	Session	Non-session
All APs	480/ 665	566/ 1017	408/ 372
Type A	228/ 327	419/ 598	69/ 102
Type B	252/ 337	147/ 419	339/ 269

(a) 11a traffic (uplink/ downlink)

AP \ Time	All	Session	Non-session
All APs	1257/ 3960	1938/ 6740	690/ 1647
Type A	734/ 2443	1473/ 4871	119/ 424
Type B	523/ 1516	465/ 1869	571/ 1222

(b) 11b traffic (uplink/ downlink)

One of possible reasons is the shortage of 11a's coverage. 5GHz WLAN's coverage is shorter than that of 11b's. Accordingly, there could be more service coverage holes in the conference area. In such a case, it is likely that 11a clients cannot sustain their connection while roaming across the conference area. The summary of the session time statistics is as follows.

- The sojourn times of 11a and 11b are similar. About 70% of the clients keep the connection with one AP less than 10 minutes.
- In the case of 11b (11a), about 90% (90%) of the clients experience fewer than 5 (2) hand-offs, and over 50% (50%) of the clients do not roam at all. About 75% (85%) of the clients keep the session for less than 1 hour.

From Figs. 8, 9, and 10, we can analyze 11b and 11a sojourn times in detail. Actually, the sojourn times can be divided into two groups: one is the Short Sojourn Time (SST) group, which represents the clients which do not stay at a fixed position and pass through AP's service areas. The other group is the Long Sojourn Time (LST) group, which means the clients which associate with an AP, and stay at the fixed location (or more exactly within the AP's service area) for a while. Even though 11a and 11b sojourn times are similar, the SSTs of 11a and 11b can be different. We assume 100 seconds as the threshold value between SST and LST.

Fig. 8 shows the conditional CDF of SST given that clients do not roam across APs. In the case of 11b, very short sojourn times less than 10 seconds occupy a major portion. The reason for that can be illustrated by client 3 in Fig. 10. When an 11b client is near the boundary of an 11b AP's coverage, the client can associate with the AP. However, the radio signal may not be strong enough to sustain the connection, and hence the client repeats to lose and recover its connection. In such a case, the client sojourn time can be very short. In Cisco AP, when AP sends packets to a client and cannot receive the ACK of those packets, it logs "MAXRETRIES" event and send a disassociation message to the client in order to disconnect session.

The 11a case could be similar to the above-described 11b case when the client is near the edge of an AP's coverage. However, such a behavior is not visible in the case of 11a since the distribution of 11a's SST is relatively dense around the time of 60 seconds as explained below. This seems to be due to the disassociation time-out events, where the default time-out threshold is 60 seconds with the Cisco APs. When a client goes outside an AP's service coverage without an explicit disassociation with its AP like client 2 in Fig. 10, the AP waits for the client to be active again for a time-out threshold. Upon a time-out, the AP disassociates the client from itself by transmitting a disassociation frame to the client, and deleting the association information. In the 11a's case, the service coverage of the APs is relatively smaller than that of 11b's, and hence an 11a client is more likely to fail to roam to another AP when user roams around. That is the reason why there is a relatively bigger jump near 60 seconds in the CDF of SST.

Fig. 9 shows the CDFs of SST of 11a and 11b given that clients in consideration perform hand-offs at least once. We can again observe that 11a's SST is longer than that of 11b in average. One of the possible reasons is also the service coverage difference between 11a and 11b. In Fig. 10, the client 1 goes from AP 1's coverage to AP 2's. The 11b APs' combined coverage is seamless, and hence an 11b client can sustain its connection during the movement, and reassociates with AP 2 when the client approaches near point A. However, an 11a client must wait until it reaches AP 2's coverage, and can reassociate with AP 2 around point B. Therefore, the 11a's SST is longer than that of 11b as was observed in Fig. 9, although the difference is not so huge. From Fig. 8 and 9, we can conclude that the 11a's SST is generally longer than 11b's due mainly to the service coverage difference.

As far as we understand, the session time durations can be useful to determine the DHCP lease time [7]. For public WLAN service [9], to determine the optimal lease time is helpful to reuse a pool of IP addresses efficiently for many clients.

D. User Mobility

Here, we present further analysis on the user mobility from the hand-off event measurements. Table III shows the number of visited APs per client during the entire meeting

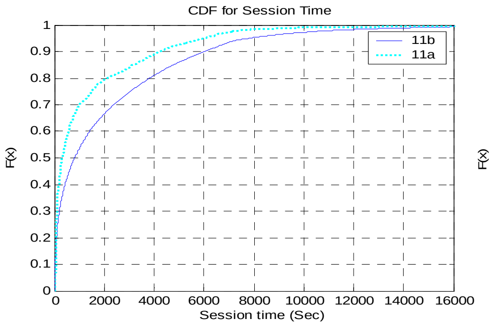

Fig. 5. The Session time duration

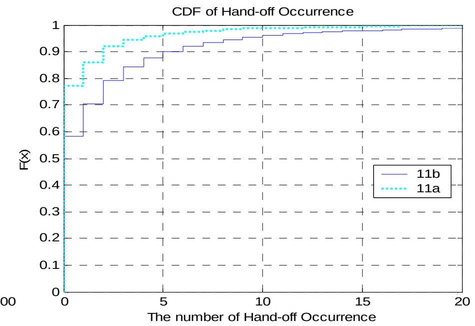

Fig. 6. The # of hand-off occurrences during one session

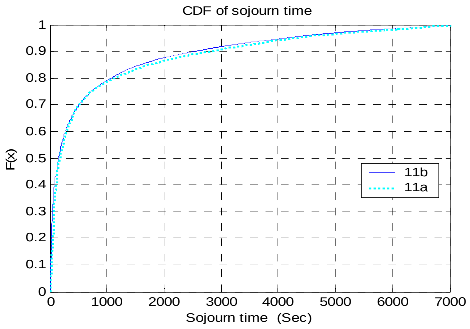

Fig. 7. The sojourn time bet. an AP and a client

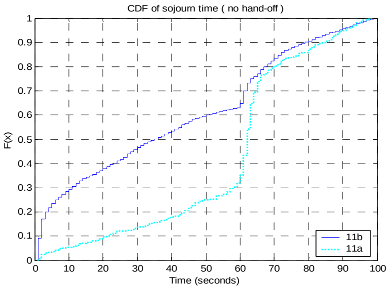

Fig. 8. The SST distribution
when no hand-off

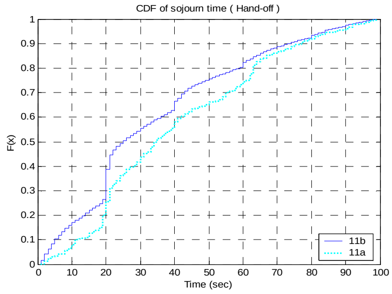

Fig. 9. The SST distribution
when hand-off situation

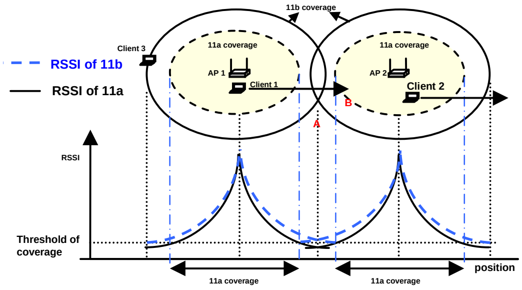

Fig. 10. Relative service coverage of 11a and 11b APs

period. The 11b users visited about 2 more APs than 11a users in average. It implies that 11a and 11b users' physical mobility is not very different as it should be.

Fig. 11 is a plot to show the number of hand-offs as a function of old and new AP pairs, where old AP means the AP which the handing-off client associated with before a hand-off, and new AP is the AP which the client reassociates with upon the hand-off. In both plots, X-axis represents the old APs while Y-axis represents the new APs. The APs are sorted by the last octet of their IP addresses, and are mapped into the axis so that the neighboring APs can be located close along the axes, because we intentionally assigned the IP address according to physical locations of APs. The left graph of Fig. 11 shows the 11b clients' hand-off pattern, and the right one shows the 11a clients'. Obviously, the hand-off of 11b clients is much more frequent than that of 11a's. As the position of X-Y plane is closer is to the diagonal line, the value of Z is larger generally. Because the geographical distance between two APs increases proportional to the distance along the axes in general, the AP pairs near the diagonal line are geographically close ones. Thus, large Z values near the diagonal line of the X-Y plane is a natural hand-off pattern. Between some AP pairs, which are located far from each other, we observe quite a few hand-off events as well. It is because the distance of two APs is close enough to directly hand off, although the X-Y plane mapping is related to the order of AP's deployment. Note that the 11b's ratio of the hand-off between near APs to that of far APs is larger than the 11a's ratio; the Z-axis values near the diagonal line are relatively larger than that of other regions in case of the 11a plot compared to the 11b plot. This phenomenon can be explained by the fact that 11a coverage is smaller than the 11b's, so the possibility of hand-off between far APs decreases in the case of 11a. This fact quite matches with the coverage hole hypothesis explained in the previous subsection.

Another good evidence for the coverage hole can be also observed in Fig. 11. The front part of 11a's diagonal region of the X-Y plane is quite small compared to the 11b's plot. The front part is actually mapped to the APs located on the first floor, i.e., those belong to type B, and are deployed in a more geographically scattered manner. On the first floor, the 11b clients can hand off without disconnection because the 11b coverage is enough to maintain the connection until the hand-off. On the other hand, 11a clients cannot roam seamlessly due to the shorter RF propagation of 11a. However, this phenomenon can be different depending on AP model and RF transmission power.

Next, we focus on asymmetric hand-off behavior as another issue on mobility. Fig. 12 shows the distribution of "asymmetric hand-off ratio," which varies for different old and new AP pairs. The definitions of old AP and new AP are the same as those of Fig. 11. "Asymmetric Hand-off Ratio" AHR(A,B) between AP A and AP B is defined as follows:

$$AHR(A,B) = \begin{cases} \dfrac{|HO(A,B) - HO(B,A)|}{HO(A,B) + HO(B,A)}, & HO(A,B) + HO(B,A) > T \\ 0, & \text{otherwise} \end{cases}$$

where HO(A,B) is the number of hand-off events from AP A to AP B, and T is a threshold value. If the number of hand-offs between AP A and AP B is very small, this ratio can be exaggerated. For example, if the number of hand-offs from A to B is 1 and that for the opposite direction is 2, AHR(A,B) is 1/3. It is a quite large value even though the difference of hand-off numbers is only 1. Therefore, AHR can be meaningful only when the hand-off number is large enough. We set 100 as the threshold value in the case of 11b, and 10 in the case of 11a. Generally 11b's hand-off is more

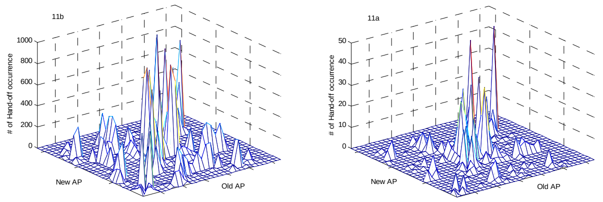

Fig. 11. The distribution of hand-off per an AP pair (left: 11b, right: 11a)

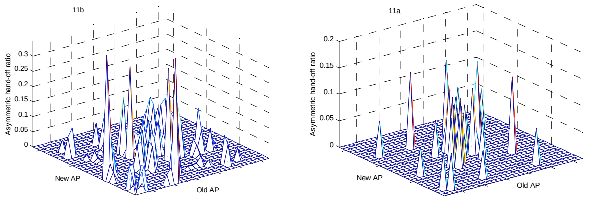

Fig. 12. The distribution of symmetric hand-off ratio (left: 11b, right: 11a)

TABLE III.
The statistics of the number of visited APs

Values	Mean	Median	Maximum
11b	10.8	11	28
11a	8.8	8	21

asymmetric than 11a's, because 11b's hand-off occurs in more various situations than 11a's case. However, except a few cases, the hand-off asymmetry is not severe. Asymmetry can be caused by topological deployment, for example, "L" shaped aisle and inter-floor situation [4]. When a client goes out of a certain AP's coverage, the client starts searching neighboring APs. Usually, the client tends to keep its association until the RF signal decreases under a threshold, even though the signal strength from other APs is greater than that of the current AP. Therefore, when a client moves from some place to another place and vice versa, where the client starts to roam is different depending on the moving direction. If the place is not an open space, especially, a dense area with APs or a building with a complicated structure, whenever the client searches APs, the scanning result can be quite different, so the client moves to different APs which it will not follow when it moves to the opposite direction. This asymmetry can effect on the performance to predict the future APs for reducing hand-off latency.

V. CONCLUSION

In this paper, we have presented our analysis of network traffic and user mobility patterns of a conference type network based on the IETF meeting. It is the first result of 11a measurement of public environment. Our measurement includes 1293 clients, 32 APs for 5 days. From the measurement data, we can summarize the result as follows.

- User arrival pattern: the user arrival pattern is highly correlated with the meeting schedule. We divide the deployed APs into two groups. Type A represents the APs in the large meeting rooms, and type B is for the rest of the APs. Type-A APs show a large variation in the associated clients between session times and non-session times. The number of users associated with type-A APs during session time is about 20 times as many as one of non-session time. In case of type B, the number is only doubled when a session is opening. When a session is not held, type-B APs are more crowded by clients than type-A APs. The tendency is about the same for both 11a and 11b.
- Traffic pattern: the traffic pattern is similar to the user arrival pattern, which highly depends on the meeting schedule. Downlink traffic is heavier than uplink traffic, and total 11b traffic is much more than 11a traffic, because the number of 11b users is larger than that of 11a. However, 11a traffic per client is heavier than that of 11b, by 3.3 times. When we design 11b/11a dual-mode networks, we should assume the traffic per user and channel capacity differently for 11a and 11b.
- Session pattern: the sojourn times of 11a and 11b are similar. About 70% of the clients keep the connection with one AP for less than 10 minutes. During a session, over a half of the clients do not change their AP (i.e., no hand-off). 11b clients perform hand-off more frequently than 11a clients; while 90% of 11b clients hand off up to 5 times, 90% of 11a clients hand off for only up to 2 times. One of possible reasons is more service coverage holes of the 11a network compared to the 11b network.
- User mobility: hand-offs between geographically-distant APs are more frequent in the 11b network than in the 11a network because 11a wireless signal cannot reach farther than 11b. Hand-off asymmetry is an important issue for predicting future APs, but the asymmetry phenomenon is not often observed.

In the near future, many companies, and hotspots will be provided 11a/11b (or 11g) dual (or triple)-mode networks. We expect that our analysis on 11a/11b user behavior and usage pattern will be helpful for 11a/b dual-mode (or 11a/b/g triple-mode) WLAN network designers in the future.

REFERENCES

[1] D. Tang and M. Baker, "Analysis of a Local-Area Wireless Network," in *Proc. ACM MOBICOM'00*, Aug. 2000.

[2] D. Kotz and K. Essien, "Analysis of a Campus-Wide Wireless Network," in *Proc. ACM MOBICOM'02*, Sep. 2002.

[3] F. Chinchilla, M. Lindsey, and M. Papadopouli, "Analysis of Wireless Information Locality and Association Patterns in a Campus," in *Proc. IEEE INFOCOM'04*, Mar. 2004.

[4] A. Mishra, M. Shin, and W. A. Arbaugh, "Context Caching Using Neighbor Graphs for Fast Handoffs in a Wireless Network," in *Proc. IEEE INFOCOM'04*, Mar. 2004.

[5] M. Shin, A. Mishra, and W. A. Arbaugh, "Improving the Latency of 802.11 Hand-Offs using Neighbor Graphs," in *Proc. ACM MobiSys'04*, June 2004.

[6] Cisco AP event and error message, http://www.cisco.com/univercd/cc/td/doc/product/wireless/aero1400/br1410/brscg/p11err.htm, online link

[7] A. Balachandran *et al.*, "Characterizing User Behavior and Network Performance in a Public Wireless LAN," in *Proc. ACM SIGMETRICS'02*, June 2002.

[8] D. Kim, S. Kim, H. Jung, and S. Gweon, "WLAN Interference Reduction Method by Channel Assignment with Field Test Results," in *Proc. JCCI'03*, April 2003.

[9] Y. Choi, J. Paek, S. Choi, G. Lee, J. Lee, and H. Jung, "Enhancement of a WLAN-Based Internet Service in Korea," in *Proc. ACM WMASH'03*, San Diego, USA, Sep. 19, 2003.